



Data Virtualization with SharePoint 2013

using

Enterprise Enabler[®]

*A full-spectrum integration platform
fine-tuned specifically for SharePoint*

Copyright © 2014 Stone Bond Technologies, L.P. All rights reserved.

The information contained in this document represents the current view of Stone Bond Technologies on the issue discussed as of the date of publication. Product names mentioned may be trademarks of their respective companies.

This white paper is for information purposes only.

Stone Bond Technologies may have patents, patent applications, trademark, copyright or other intellectual property rights covering the subject matter of this document.

Except as expressly provided in any written license agreement from Stone Bond Technologies, the furnishing of this document does not give you any license to these patents, trademarks, copyrights or intellectual property.

Stone Bond Technologies, L.P.

1021 Main Street Suite 1550

Houston, TX 77002

713-622-8798

www.stonebond.com

SharePoint-Centric Data Virtualization

What is Data Virtualization?

It is making data from multiple, usually dissimilar, sources available for use as if it were a single source, to be acted on using the full range of reading and writing back to the original sources, without requiring a formal data model or staging database. In current vernacular, data virtualization assumes use is only on-demand, not physically moved.

Enterprise Enabler's data virtualization definitions can also be used for ETL, where data is federated as it is being physically moved to some destination.

Why is Data Virtualization Important to SharePoint?

1 Most SharePoint apps need data from multiple sources.

2 Users and apps interact with live data as if it were a single source.

3 Eliminate staging databases.

4 Turn dashboards into interactive consoles.

The Synergy of SharePoint and Data Virtualization

It's a powerful combination: Data Virtualization and SharePoint. The most recent releases of SharePoint, 2010 and 2013, support broader-than-ever opportunities for data integration and data-centric implementations. SharePoint has evolved to include its availability as a service in the cloud and a hybrid architecture, along with its, fully on premise model.

The availability of these alternate architectures is a compelling argument for adopting SharePoint for many types of applications beyond the historical content management uses. The rich Business Connectivity Services (BCS) provide a range of features that are data-centric, elevating SharePoint to another level as a platform for applications, dashboards, and BI tools.

As implementations mature, though, certain requirements are coming to the forefront. In particular, the reality of most applications is that data needs are not limited to a single data source, or even a single data source type, which is a limitation of not only SharePoint, but of the third-party vendors that provide connectivity. The potential of the built-in powerful BCS is not fully leveraged, and the end result is either limited scope of applications or the need to build a staging database either on premise or in the cloud, the latter of which brings additional security concerns for many customers.

The ideal use of SharePoint leverages BCS to access data live from the original sources, validates the data, presents it to the end user, or application for action or editing, and pushes changes back to the original sources without ever saving it anywhere. End user security ensures that end users only see or edit data to which they have access.

This is what Data Virtualization (DV) brings to the table, and Stone Bond's Enterprise Enabler is a full-scope integration platform, with data federation at its core. The single Integrated Development Environment (IDE) for development, testing, deployment and monitoring, is built on and fully extensible with Microsoft .Net.

SharePoint and Enterprise Enabler

Data Virtualization is not just for Business Intelligence these days. It is behind the scenes of all kinds of applications, and SharePoint is a useful platform on which to build a wide range of applications and dashboards, particularly when combined with Enterprise Enabler's powerful data federation and data virtualization capabilities. Enterprise Enabler solves the ubiquitous challenge of interacting with many data sources of various types that view and treat data in completely different ways. Its underlying integration technology ensures that information is properly aligned, validated and made available in whatever form is most appropriate for the endpoint or consuming application, without incurring a huge overhead of custom programming to configure and maintain the solution.

While Enterprise Enabler's federated data can be executed in an ETL manner to traditional SharePoint Lists or TermSets in event driven processes, it can also be available on-demand with full CRUD (Create, Read, Update, and Delete) and query-able via any data access that the SharePoint application requires. That could be Enterprise Enabler CustomConnector, OData services, or for custom SharePoint applications, ADO.Net, ODBC, SOAP and REST web services. This means that SharePoint application developers can easily incorporate data from virtually anywhere in whatever mode they already have in place, or are planning to use, for data interaction, all with secure bi-directional data access.

Cloud applications generally require the data to be in the cloud. For a many businesses, this is still considered high risk for most of their data, not only by having it reside in the cloud, but also simply because a copy must be made. For many cloud based applications the data required originates in multiple places. The traditional solution is to design a data base that will hold all the anticipated data, and then program integration utilities to update or synchronize the data periodically from the sources. It is common wisdom that this exercise consumes at best half of the development time and budget.

Using Enterprise Enabler's Data Virtualization, the cloud application can access all of the back-end data directly, as if it were in a combined database.

Whether the SharePoint application resides in the cloud or on premise, this means the data is fresh, write-backs are immediately posted, and all of the overhead of designing and maintaining the for-purpose database is eliminated. Data Virtualizations can be configured and deployed as an Enterprise Master Service (EMS) in matter of minutes in any or all of the above mentioned data access as well as JDBC for continuity with Java based applications.

Additionally, if it turns out that the decision is to move or copy the data to the cloud, Enterprise Enabler can use the same Data Federation logic configured for EMS as the basis of an ETL for physically updating or synchronizing the operational database, wherever it resides.

How Has this Synergy Been Used?

Utility Company's HR Portal

SharePoint is being selected more than ever as the platform of choice for intranet and outward-facing portals. This HR portal was designed to streamline the access, modification, and dissemination of personnel information. Employees can easily access their own information as well as policies and guidelines and such whenever needed, and management can see and enter performance and other information. Enterprise Enabler accesses data from SAP, from DB2 and other custom databases and applications. Rather than designing a staging data store to be updated periodically, the lead architect decided to use a data virtualization approach, accessing data from all the sources live, and presenting the appropriate information in SharePoint using BCS supported Enterprise Enabler (EE) CustomConnector, with Enterprise Enabler doing the heavy lifting for the data federation and write-back. Certain fields such as employee's address were configured to be editable so that the user could correct or change the address. The integration for this project was configured within three weeks, including the security pass-through using SSS.

Multi-Refinery Planning Process

This was a BPM project of sorts, providing user interfaces for separate roles of Planning, Marketing, and Operations. The "market driven manufacturing" objective faced considerable challenges in accessing data all kinds of varied applications and data sources. Logical models needed to be aligned with physical models and real-time data with planning guidelines over time frames on the order of months. Every source had different specification, units of measure, and meanings. Enterprise Enabler brought all the information together, automatically staging some for performance availability and executing complex data transformations across multiple sources. A SharePoint application served as the user interfaces, capturing data and moving it behind the scenes to specialized optimization models as necessary. Enterprise Enabler's ability to monitor schemas for change eliminated the typical data inconsistencies that are a part of systems in dynamic operations.

Leveraging SharePoint Features without Custom Coding

Business Connectivity Services in SharePoint 2010 and now including OData in 2013, along with CustomConnector, offer a fertile ground for bi-directional data virtualization. The native ability to send instructions for Create, Read, Update, and Delete (CRUD) and query, automatically elevates the potential of any application to be dramatically more interactive. Instead of dashboards that display data that has been staged in some database, a user can see and interact with live data from backend systems. This means turning dashboards or BI tools into interactive business consoles carrying updates and decisions back to the enterprise.

SharePoint's hooks to interact with data are far ahead of the generally available integration tools that carry out the data instructions. Unfortunately, most integration products only handle READ; some handle READ and QUERY; most handle only a small set of endpoint types; most handle only a single source at

a time; most require considerable custom coding to set up integrations and data transformation. While BCS features connectivity including OData, CustomConnectors, and a few select applications and databases, it does not have any facility for data federation across disparate sources. Even SSIS and SQL Server 2014 simply cannot do bi-directional data virtualization, even with a considerable amount of coding.

Stone Bond’s Enterprise Enabler covers all of these features without coding, including data virtualization, which seamlessly aligns and combines from any number of different sources. It is the only enterprise grade integration platform based on .Net technology, making it a natural clean extension for SharePoint or any .Net based application.

Enterprise Enabler and SharePoint Combined Architectures

Configurations of SharePoint with Enterprise Enabler

	SharePoint Online	Private Cloud	Hybrid	On- Premise
SharePoint Client		SP Online or Private		OP
SharePoint Server/Farm		Private	OP	OP
Enterprise Enabler® Server				OP
EE Custom Connector	None	On SharePoint Server	None	On SharePoint Server
Endpoint Locations	Anywhere	Anywhere	Anywhere	Anywhere
Delivery Mode	OData Only	All Modes	OData Only	All Modes

OP = (On-Premise) **MODES** = OData, SOAP, REST, ODBC, ADO.Net, EE CustomConnector

SharePoint 2013 literature and technical notes discuss three distinct architectures, Cloud-Only, Hybrid, and On-Premise. Each of these can take advantage of Enterprise Enabler’s data Virtualization technology to quickly configure full CRUD/Query access to federated data sets combining sources in the cloud and on-premise. Enterprise Enabler’s architecture is flexible in that it can accommodate various

Key features of Enterprise Enabler

1. Connects to anything using AppComms™, intelligent endpoint connectors

2. Handles all integration with data federation for all patterns: ETL, EAI, SOA, DV

3. Native Transformation Engine that interacts simultaneously with multiple endpoints and performs complex transformations and alignment of live data from the endpoints

4. Data Virtualization with write-back that manages security and transaction roll-back across multiple sources.

5. Powerful Data Workflow engine to manage potentially complex combinations of data flows, data cleaning and validation, error handling logic, cache set-up and update and other processes and operations.

6. Pro-active monitoring for change in schemas – the guardian of stability and agility.

7. Single Integrated Development Environment (IDE) to design, build, test, deploy and monitor integrations.

configurations, with or without the EE Custom Connector, EE Server located in the cloud or on-premise, and handling security appropriately according to the selected SharePoint architecture chosen.

• **SharePoint Online, Cloud Only**

- EE OData services on-premise that are exposed to the internet can be used (with some security).
- EE could reside in the cloud as an OData service for BCS.

• **Private Cloud, a private cloud is owned and operated by a company and could have a vpn channel to the on-premise domain or not.**

- With secure VPN channel to on-premise network/domain
 - CustomConnector for EMS can be installed that connects to on-premise EE Server.
 - EE OData services for EMS can be used.
- Without secure VPN channel to on-premise network/domain
 - EE OData services for EMS can be used by exposing the services to internet (with security).

• **Hybrid**

- When SharePoint online uses inbound connection to on-premise SharePoint 2013 through reverse proxy.
 - BCS OData connector sends requests through the reverse proxy to on-premise SharePoint 2013 BCS which forwards the request to OData service endpoint.
 - EE OData services for EMS are used in this scenario.
- In order to honor security in this model, EE should be on Premise with a SharePoint Server to allow the reverse proxy trust.
- On-premise SharePoint 2013 farm can install CustomConnector for EMS (for on-premise use).

• **On-Premise**

- CustomConnector for EMS can be used.
- EE OData services for EMS can be used.

Additional Value from Enterprise Enabler

- **Patterns and Caching.** Because of the range of integration modes incorporated into Enterprise Enabler, compound and complex data patterns can easily be configured. This attribute of Enterprise Enabler allows unfettered creativity for architects to handle whatever situation is at hand. For example, suppose the objective is to provide a customer portal that uses data from Oracle, SAP, and Salesforce.com. Each time a customer refreshes his screen, each backend system is hit. This could cause performance issues and potentially impact the availability of the backend systems. The architect can elect to cache some of the data that is known to be relatively constant. He can set up an in-memory or in-database cache and quickly map the integration to update the cache in an ETL manner, with data workflow defining conditions when the update occurs.
- **Data Masters.** Each time an Enterprise Master Service (EMS) is defined, it becomes a candidate to become a sanctioned Data Master, which is a data entity that not only defines that data set for a particular set of uses, but also incorporates all of the data sources and all the alignment and transformation rules. The EMS can be produced in every mode: EE BCS CustomConnector, ODBC, ADO.Net, JDBC, REST, SOAP, OData, others, and ETL.
- **Integration Integrity.** Enterprise Enabler's patented process, IIM, monitors schemas for change, performs impact analysis, and sends notifications, of potential issues. IIM is used at development time as well as in production.
- **ROI.** The discussion would not be complete without a comment on ROI. It is important to reflect on the considerable savings that come from using the Enterprise Enabler platform. No specialized skillsets are required, most technical work is performed behind the scenes by Enterprise Enabler as it constructs and executes metadata instructions. Not staging databases need to be designed, fed, and maintained.

Enterprise Enabler Developer Experience

One of the most important attributes of Enterprise Enabler is that it packs ETL, EAI, ESB, and Data Virtualization into a single product with a single User Interface. The design, testing, deployment and monitoring of the solution is performed without ever leaving this environment. Here is an example of what a .Net developer or a DBA would do to configure a SharePoint connectivity that combines data from, for example, Salesforce.com, SAP, and a SQL Server database.

1. **Log in to Enterprise Enabler** so it can maintain its audit trail, and can control what you see and are able to do within the product.
2. **Build and save a Template for each source** to indicate which data from which specific instance of the application you are interested in. Select the AppComm (intelligent connectivity) for the particular type of source, e.g., SAP. Enter the credentials or connection string. You will be

presented with a selection mechanism appropriate to the application so you can select the data you want. For Salesforce you will see all the objects/tables in your instance, including the ones you have custom configured. Pick the ones you need.

3. **Build a virtual destination Template.** Since the destination doesn't exist yet, you may add fields and definitions one by one or reuse parts of existing templates to define what you want the entity to include when you use it in SharePoint.
4. **Open the Active Integration Interface** and select the source Templates you just built, or reuse some that were already configured. Select your virtual destination Template. You will see a three-pane screen, with your sources, Salesforce, SAP, and SQL Server on the left and destination on the right.
5. **Define Relationships across the sources** by relating key fields, just as you would build a relationship in a database.
6. **Map the sources to the destination.** Drag, drop, add functions, filters, validations, web service calls, using the rich mapping features. If you need to add transformations that are not available, just call the function builder, use the intellisense of the C# or VB.Net embedded editor to write it, and compile and test it within the same User Interface.
7. **Test the Mapping.** From that same screen, simply click the Run Once button. EE executes your map live against the backend systems and presents you with the resultant data for your review and validation. If you see something that's not right fix it and test again Notice you do not need to do leave the environment or even the mapping screen. Save the map.
8. **Make one or more Enterprise Master Services (EMS)** by selecting which modes you want packaged, and your data virtualization will be automatically generated and ready to be selected for association from SharePoint as you build your application.

There is no further scripting or cajoling required to get the federated entity working. In the SharePoint application, CRUD functions are handled by Enterprise Enabler, passing security tokens back to the source systems, to ensure proper authority when users interact with your application.

Summary

Enterprise Enabler is the only data integration product that can take full advantage of SharePoint 2010 and 2013 BCS feature sets without significant custom programming work. Even Microsoft tools like SSIS and SQL Server 2014 require tremendous programming efforts and bi-directional data virtualization is simply not possible. The Data Virtualization Enterprise Master Services interface directly with SharePoint in all modes and architectures, bringing the benefit of live data simultaneously from multiple disparate sources for reading as well as updating.

The ROI, coupled with ease of implementation and maintenance, along with the richness of the platform offer a compelling and unparalleled opportunity to simplify your business and IT Infrastructure.

Microsoft and Stone Bond Collaboration

In late 2008 Stone Bond began working with Microsoft's SharePoint Development team to build connectivity, using Enterprise Enabler, to pre-alpha SharePoint 2010 BDC (now BCS, External Content). In the absence of any Microsoft tools to build or test full-featured connectivity, the Microsoft team worked closely with Stone Bond to define and test the complex features. At the dev kitchen in early 2009, Stone Bond demonstrated in fifteen minutes, configuring Enterprise Enabler to federate live data from SQL Server and Salesforce, and displaying it in SharePoint using the BDC external list, including writing back the end user modified data to Salesforce. At the 2010 Airlift conference, Stone Bond was the only non-Microsoft presenter to demonstrate features of the then new SharePoint. Microsoft has since incorporated some of the capabilities Stone Bond developed and demonstrated into their own 2013 release, but still, without Enterprise Enabler, falling short of providing the variety of endpoints, of the data federation and virtualization as well as ease and speed of configuration.

About Stone Bond

Stone Bond Technologies, L.P. is the provider of Enterprise Enabler®, the leading business integration software for data virtualization, federation and orchestration delivering the industry's fastest time-to-value. As the only integration solution to scale both up and down to meet business needs, Stone Bond provides software solutions that help businesses improve their bottom line by streamlining interaction among business systems. To learn more about Stone Bond Technologies please visit, www.stonebond.com.

See more at www.stonebond.com